

IN THE CLAIMS

1. (Currently Amended) A method for providing a file system snapshot, comprising:
  - creating an initial snapshot dataset for a source file having at least one of metadata and data in a file system, wherein the initial snapshot dataset contains substantially no data and no metadata;
  - accepting, subsequent to the creating, a command to modify metadata in a second inode within the source file;
  - copying, in response to accepting the command to modify metadata, at least a portion of the metadata within the second inode into a first inode within the snapshot dataset; and
  - storing, in response to the accepting, disk address values into a disk address of the first inode that are equal to a ditto address, the ditto address indicating that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot.
2. (Previously Presented) The method of claim 1, further comprising:
  - appending data to the source file;
  - copying, in response to the appending, to the first inode in the snapshot dataset at least a portion of metadata within the second inode; and
  - storing, in response to the appending, disk address values into the first inode equal to the ditto address to indicate that the disk address is stored in one of an inode of the file system and a subsequent snapshot.
3. (Previously Presented) The method of claim 2, further comprising:
  - modifying a source file by one of either overwriting and deleting one of the data blocks corresponding to the source file; and
  - copying, in response to the modifying, to the first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, wherein the first inode includes a disk address of the data block which was written in the snapshot dataset.

4. (Previously Presented) The method of claim 3, wherein the ditto address is an invalid disk address, the method further comprising:
  - accessing the first inode of the snapshot dataset corresponding to the source file;
  - determining whether the first inode includes a valid disk address;
  - reading, in response to determining that the first inode includes a valid disk address, a data block referenced by the disk address; and
  - retrieving, in response to determining that the first inode includes the ditto address, the second inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.
5. (Previously Presented) The method of claim 3, further comprising:
  - copying, in response to the modifying, to the first inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, wherein the first inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.
6. (Previously Presented) The method of claim 5, wherein the ditto address is an invalid disk address, the method, further comprising:
  - accessing the first inode corresponding to the source file;
  - determining whether the first inode includes a valid disk address;
  - retrieving, in response to determining that the first inode includes a valid disk address, an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block; and
  - retrieving, in response to determining that the first inode does not include a valid disk address, the second inode of the source file, then retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at

least one data block referenced by at least one disk address in the indirect block.

7. (Currently Amended) A system for providing a file system snapshot, comprising:
  - means for creating an initial snapshot dataset for a source file having at least one of metadata and data in a file system, wherein the initial snapshot dataset contains substantially no data and no metadata;
  - means for accepting, subsequent to the creating, a command to modify metadata in a second inode within the source file;
  - means for copying, in response to accepting the command to modify metadata, at least a portion of the metadata within the second inode into a first inode within the snapshot dataset; and
  - means for storing, in response to accepting a command to modify metadata, disk address values into a disk address of the first inode that are equal to a ditto address, the ditto address indicating that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot.
8. (Previously Presented) The system of claim 7, further comprising:
  - means for appending data to the source file;
  - means for copying, in response to appending data to the source file, to the first inode in the snapshot dataset at least a portion of metadata within the second inode; and
  - storing, in response to appending data to the source file, disk address values into the first inode equal to a the ditto address to indicate that the disk address is stored in one of an inode of the file system and a subsequent snapshot.
9. (Previously Presented) The system of claim 8, further comprising:
  - means for modifying a source file by one of either overwriting and deleting one of the data blocks corresponding to the source file; and
  - means for copying, in response to modifying the source file, to the first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, wherein the first

inode includes a disk address of the data block which was written in the snapshot dataset.

10. (Previously Presented) The system of claim 9, wherein the ditto address is an invalid disk address, the system further comprising:

means for accessing a first inode of the snapshot dataset corresponding to the source file;

means for determining whether the first inode includes a valid disk address,

means for reading, in response to a determination that the first inode contains a valid address, a data block referenced by the valid disk address; and

means for retrieving, in response to a determination that the first inode does not contain a valid address, an inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.

11. (Previously Presented) The system of claim 9, further comprising:

means for copying, in response to modifying the source file, to the first inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, wherein the first inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.

12. (Previously Presented) The system of claim 11, wherein the ditto address is an invalid disk address, the system further comprising:

means for accessing a first inode corresponding to the a source file;

means for determining whether the first inode includes a valid disk address,

means for retrieving, in response to determining that the first inode includes a valid disk address, an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block; and

means for retrieving, in response to determining that the first inode does not

include a valid disk address, the second inode of the source file, retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block.

13. (Currently Amended) A computer readable medium including computer instructions for providing a file system snapshot, the computer instructions comprising instructions for:

creating an initial snapshot dataset for a source file having at least one of metadata and data in a file system, wherein the initial snapshot dataset contains substantially no data and no metadata;

accepting, subsequent to the creating, a command to modify metadata in a second inode within the source file;

copying, in response to accepting the command to modify metadata, at least a portion of the metadata within the second inode into a first inode within the snapshot dataset; and

storing, in response to the accepting, disk address values into a disk address of the first inode that are equal to a ditto address, the ditto address indicating that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot.

14. (Previously Presented) The computer readable medium of claim 13, the computer instructions further comprising instructions for:

appending data to the source file;

copying, in response to the appending, to the first inode in the snapshot dataset at least a portion of metadata within the second inode; and

storing, in response to the appending, disk address values into the first inode equal to the ditto address to indicate that the disk address is stored in one of an inode of the file system and a subsequent snapshot.

15. (Previously Presented) The computer readable medium of claim 14, the computer instructions further comprising instructions for:

modifying a source file by one of either overwriting and deleting one of the data blocks corresponding to the source file; and

copying, in response to the modifying, to the first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, wherein the first inode includes a disk address of the data block which was written in the snapshot dataset.

16. (Previously Presented) The computer readable medium of claim 15, wherein the ditto address is an invalid disk address, the computer instructions further comprising instructions for:

accessing the first inode of the snapshot dataset corresponding to the source file;

determining whether the first inode includes a valid disk address;

reading, in response to determining that the first inode includes a valid disk address, a data block referenced by the disk address; and

retrieving, in response to determining that the first inode includes the ditto address, the second inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.

17. (Previously Presented) The computer readable medium of claim 15, the computer instructions further comprising instructions for:

copying, in response to the modifying, to the first inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, wherein the first inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.

18. (Previously Presented) The computer readable medium of claim 17, wherein the ditto address is an invalid disk address, the computer instructions further comprising

instructions for:

accessing the first inode corresponding to the source file;  
determining whether the first inode includes a valid disk address,  
retrieving, in response to determining that the first inode includes a valid disk address, an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block; and  
retrieving, in response to determining that the first inode does not include a valid disk address, the second inode of the source file, then retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block.

19. (Currently Amended) A system for providing a file system snapshot, comprising:  
an initial snapshot dataset for a source file containing data in a file system,  
wherein the snapshot dataset is substantially empty;

means for accepting a command to modify metadata in a second inode within the source file; and

a first inode in the snapshot dataset, the first inode comprising metadata copied from second inode corresponding to the source file, wherein the first inode is generated in response to accepting the command and wherein a ditto address is stored in a disk address of the first inode, the ditto address indicating that the stored disk address is an invalid disk address and indicates that the disk address is an invalid disk address and also indicates that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot.

20. (Previously Presented) The system of claim 19, wherein in the first inode, the metadata from the first inode is copied from the second inode corresponding to the source file, wherein the first inode is generated only when the data block corresponding to the source file is appended and wherein the ditto address is inserted into the first inode.

21. (Previously Presented) The system of claim 20, further comprising:

a data block corresponding to the source file in the snapshot dataset, wherein the data block is copied to the snapshot dataset when the original data block is overwritten; and

a first inode in the snapshot dataset, the first inode containing metadata copied from an inode in the source file, wherein the first inode is generated when the data block corresponding to the source file is overwritten or deleted and wherein the first inode includes a disk address of the data block which was written in the snapshot dataset.

22. (Previously Presented) The system of claim 21, further comprising:

a first inode in a snapshot dataset, the first inode corresponding to a data block within a source file;

a ditto address value stored in the first inode to indicate an invalid disk address; and

an inode of the source file referencing the data block.

23. (Previously Presented) The system of claim 21, further comprising:

a first inode in a snapshot dataset, the first inode corresponding to an indirect block within a source file;

a ditto address value stored in the first inode to indicate an invalid disk address and to indicate that the disk address is an invalid disk address and that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot; and

an inode of the source file referencing the indirect block.

24. (Currently Amended) A method for deleting a first snapshot of a file system, comprising:

determining the existence of a snapshot that is older than a first snapshot;

determining, in response to determining that there is an older snapshot, the existence of a ditto address in a disk address of an inode of the older snapshot to an inode or a data block in the first snapshot, wherein the ditto address indicates an invalid

disk address and that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot; and

deleting, in response to determining that there is no older snapshot, any inode or data block in the first snapshot.

25. (Previously Presented) The method of claim 24, further comprising:

wherein if there is a ditto address in the older snapshot, copying to the older snapshot the metadata in an inode or data block of an inode in the first snapshot and deleting any inode or data block in the first snapshot; and

wherein if there is no ditto address in the older snapshot, deleting any inode or data block in the first snapshot.

26. (Currently Amended) A method for restoring a first snapshot of a file system, comprising:

accepting a request to read data from a first snapshot

determining if there is a most recent snapshot that is not the first snapshot;

copying, in response to accepting the request and in response to determining that there is a most recent snapshot that is not the first snapshot, to the most recent snapshot any inode or data block in the file system referenced by the most recent snapshot, which shall be modified by the restoration of the first snapshot;

wherein if there is an inode or a data block in the first snapshot, copying the inode or data block in the first snapshot to the file system

determining that there is a ditto address in a disk address of an inode of the first snapshot wherein the ditto address indicates an invalid disk address and also indicates that the true disk address for the actual data block is stored in one of an inode of the file system and a subsequent snapshot; and

copying, in response accepting the request and in response to determining that there is a ditto disk address in the first snapshot, wherein the ditto address indicates an invalid disk address, to the filesystem the inode or data block of the most recent snapshot that corresponds to an inode with the ditto disk address and that contains a valid disk address.